

The body in mind

Understanding cognitive processes

Mark Rowlands

University College, Cork



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK <http://www.cup.cam.ac.uk>
40 West 20th Street, New York, NY 10011-4211, USA <http://www.cup.org>
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Mark Rowlands 1999

This book is in copyright. Subject to statutory exception and to the provisions
of relevant collective licensing agreements, no reproduction of any part may
take place without the written permission of Cambridge University Press.

First published 1999

Printed in the United Kingdom at the University Press, Cambridge

Typeset in 10/12pt Monotype Bembo in QuarkXPress™ [SE]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Rowlands, Mark.

The Body in Mind: Understanding Cognitive Processes / Mark Rowlands.

p. cm. – (Cambridge Studies in Philosophy)

Includes bibliographical references and index.

ISBN 0 521 65274 X (hardback)

1. Philosophy of mind. 2. Mind and body. 3. Cognition.

4. Externalism (Philosophy of mind) I. Title. II. Series.

BD418.3.R78 1999

128'.2 – dc21 98-45620 CIP

ISBN 0 521 65274 X hardback

Contents

<i>Preface</i>	<i>page ix</i>
1 Introduction: 'A picture held us captive'	1
PART I PSYCHOTECTONICS	
2 Introduction to Part I: 'Don't work hard, work smart'	21
3 Environmentalism and what it is not	31
4 Environmentalism and evolution	64
5 Perception	100
6 Memory	119
7 Thought	148
8 Language	172
PART II PSYCHOSEMANTICS	
9 Introduction to Part II: the need for and the place of a theory of representation	205
10 Two theories of representation	212
11 Environmentalism and teleological semantics	230
<i>References</i>	258
<i>Index</i>	267

Introduction: ‘A picture held us captive’

1.1 TWO PROJECTS AND A PICTURE

Any attempt to understand how minds work must address, at the very least, two questions. The first is essentially an engineering question. One way of putting the question would be: *how can one build a mind?* This project is an engineering one. And, adopting a neologism first coined by Colin McGinn (1989), I shall refer to it as the project of *psychotectonics*: ‘psycho’, here, pertaining to minds, and ‘tectonics’ deriving from the Latin verb for building. Psychotectonics, then, is the science of building minds. In order to begin the project of psychotectonics, one must first have a reasonably adequate grasp of the things a mind can do, a grasp of the various functions of the mind. Then, it is thought, one must proceed to show, firstly, how these functions can be broken down into component sub-functions and these sub-functions broken down into sub-sub-functions, and so on, and, secondly, how these progressively more and more simple functions can be realized in progressively more and more simple mechanisms. To understand how to build a mind, it is claimed, is to be able to effect this sort of functional and mechanistic decomposition. This is a standard account of what is involved in psychotectonics, an account enshrined in David Marr’s (1982) famous tripartite distinction between *computational*, *algorithmic*, and *physical* levels of analysis; whose basic idea is reflected in Dennett’s (1978b) distinction between *intentional*, *design*, and *physical* stances; whose ethos is captured in the general project, also endorsed by Dennett (1978b) among others, of *homuncular functionalism*. This, then, is a very orthodox picture of how to do psychotectonics, and, while the picture might be a little worn in places, it is, I think, broadly correct. I do not, therefore, propose to challenge it, although I shall try to show that many of its proponents work with an unduly narrow conception of what a computational specification of the mind

should look like, and, consequently, that their conception of the algorithmic and physical realizations of this specification is unrealistic. The general idea of psychotectonics as functional and mechanistic decomposition, however, I shall accept, indeed presuppose. Psychotectonics, the project of building minds, occupies Part I of this book.

In order to understand minds, it is also necessary to understand how they can do what they can do. That is, it is necessary to understand how they come to possess those features considered essential to them. In recent discussions of the mind, two of its features loom large: *consciousness* and *intentionality*. Of consciousness, I shall have nothing to say. My suspicion, for what it is worth, is that the problem of consciousness is one that needs dissolution rather than a constructive solution. *If* dissolving the problem of consciousness requires dissolving the picture of the mind upon which it is built, and *if* this book plays a role in dissolving that conception of the mind, then the arguments of this book might be considered relevant to the problem of consciousness. But that this is so is not something I shall assume. My concern, and the principal concern of Part II of this book, is with intentionality, with the aboutness or directedness of states of mind. It may ultimately turn out to be the case that one cannot understand intentionality independently of understanding consciousness, that the two are conceptually interlinked in such a way that forming an adequate conception of the one requires adequately conceptualizing the other (McGinn 1991). This may turn out to be the case, but I suspect not. And even if it does turn out to be so, there is no guarantee in advance of which will turn out to be conceptually prior. So, I shall assume that it is possible to understand intentionality, at least to some extent, without understanding consciousness. This assumption can be questioned, but it is by no means idiosyncratic. Indeed, the assumption is fairly standard.

Most recent discussions assume what Cummins (1989) calls a *representational theory of intentionality*; that is, they assume that the intentionality of mental states reduces to the representationality of mental representations. I, also, shall assume that this is the case, and, consequently, that the project of accounting for intentionality reduces to the project of accounting for how representations represent; how physical states can have semantic properties. This project I shall refer to as that of *psychosemantics*, employing another well-known neolo-

gism, but this time in a slightly different way from that intended by its author Jerry Fodor. Psychosemantics, as I shall understand it, is the project of accounting for the representationality of representations, for how representations represent, or possess semantic properties. Again, in common with most recent discussion, I shall assume that this is a reductionist project. Representation, or semanticity, is to be explained or accounted for in terms that are non-representational, non-semantic. Psychosemantics, in this sense, is the principal concern of Part II.

The projects of psychotectonics and psychosemantics are, of course, connected. The most straightforward connection is that the two accounts must cohere in that the functional/mechanistic decomposition of the mind yielded by the psychotectonic project, when combined with the reductionist, non-semantic account of representation identified by the psychosemantic project, must, together, be sufficient to add up to intentionality. Conversely, our account of representationality identified in the psychosemantic project must be consistent with what we know, or assume, to be the correct functional/mechanistic decomposition of the mind. What unites the two projects in this book, however, are not these fairly mundane connections but, rather, a conviction that both projects are hindered by a common conception of the mind. To use a phrase popularized by Wittgenstein, a picture of the mind holds us captive. And this picture prevents us from properly understanding what is required by the projects of psychotectonics and psychosemantics.

There is a view of the mind which began life as a controversial philosophical thesis and then evolved into common sense. The view is both widespread and tenacious, not only as an explicit doctrine but, more significantly, in the clandestine influence it has on explicit doctrines of the mind. The philosophical thesis from which the view is born is spelled out by Descartes, and its association with him is sufficiently robust for it to be called the *Cartesian conception*.

According to the Cartesian conception, minds are to be assimilated to the category of substance. That is, minds are objects which possess properties. Indeed, minds can, to some extent, be conceived of as relevantly similar to other bodily organs. Just as the heart circulates blood, the liver regulates metabolism, and the kidneys process waste products, the mind *thinks*. According to official Cartesian doctrine, the major difference between the mind and these other organs is that

the mind is a *non-physical* substance. The mind and brain are distinct entities, and, while the mind may receive input from the brain, and, in turn, send information back to the brain, the two are none the less distinct. The brain is a physical organ operating exclusively on mechanical principles; the mind is a non-physical organ operating according to principles of reason. And there is, Descartes thought, no prospect of deriving the former from the latter.

The Cartesian conception has been famously ridiculed as the myth of the *ghost in the machine*. And it has been Descartes' decision to make the mind ghostly (i.e., non-physical) that has drawn the principal fire from dissenters. The dissenters' case here has largely been successful, and not many philosophers or psychologists today would regard themselves as Cartesians in this sense. Ryle's expression, however, has another facet. Not only is Descartes' mind a ghost, but it is one that is *in* a machine. This was the principal source of Ryle's ire, of course. But, whereas the revolt against ghostly views of the mind has been overwhelmingly successful, criticism of the second aspect of Descartes' view has been comparatively muted. Most theorizing about the mind is now predicated on the assumption that the mind is physical; that is, that some sort of materialism is true. However, such theorizing has been, and largely still is, predicated on the view that the mind is an internal entity, i.e., located inside the skin of any organism that possesses it. The revolt against Cartesian views of the mind has been restricted to the first aspect of Descartes' view. The other aspect, Descartes' internalism, has, until recently, largely been ignored. Most forms of materialism are, thus, also forms of internalism.

Descartes' dualism and his internalism have, arguably, the same root: the rise of mechanism associated with the scientific revolution. This revolution reintroduced the classical concept of the atom in somewhat new attire as an essentially mathematical entity whose primary qualities could be precisely quantified as modes or aspects of Euclidean space. Macroscopic bodies were composed of atoms, and the generation and corruption of the former was explained in terms of the combination and recombination of the latter. Atomism is, then, mechanistic in the sense that it reduces all causal transactions to the translation, from point to point, of elementary particles, and regards the behaviour of any macroscopic body as explicable in terms of motions of the atoms that comprise it.

It is widely recognized that Descartes' dualism stems, at least in part,

from his acceptance of mechanism. The physical world, for Descartes, is governed by purely mechanical principles. He was, however, unable to conceive of how such principles could be extended to the thinking activities constitutive of the human mind. Minds, for Descartes, are essentially thinking things and, as such, governed by principles of reason. But such principles, Descartes thought, are distinct from, and not reducible to, principles of mechanical combination and association. Rationality, for Descartes, cannot be mechanized. Each mind is, thus, a small corner of a foreign field, inherently non-mechanical, hence inherently non-physical. Descartes' dualism, in this way, stemmed quite directly from his mechanism.

Of equal significance, however, is the connection between mechanism and internalism. Mechanistic atomism is, we might say, methodologically individualist. A composite body is ontologically reducible to its simple constituents. And the behaviour of a composite body is reducible to the local motions of its constituents. Thus, if we want to explain the behaviour of a macroscopic body, we need focus only on local occurrences undergone by its parts. This methodological individualism would also have some purchase on the explanation of the behaviour of human beings, since we are also, in part, physical. It is, therefore, no surprise that minds became analogously and derivatively conceived of by Descartes, and his dualist descendants, in atomistic terms. A mind, for Descartes, is essentially a *psychic monad* (Callicott 1989). Each mind is a discrete substance insulated within an alien material cladding. Just like any other atom, the mind could interact with the physical atoms of the body. But, crucially, and again just like any other atom, the essential nature of the mind was not informed by this interaction. The rational nature of the mind is taken as an independent given, and its interaction with other atoms is extrinsic to this nature. The ghosts of this conception of the mind, and the mechanistic and individualistic conception of explanation that underwrites it, are very much with us today.

These ghosts occupy a house with many mansions. Ontological theses are entangled with epistemological ones, each giving support and succour to the other. It is genuinely unclear if any particular thesis precedes any of the others. It is more realistic to suppose, perhaps, that ontological and epistemological aspects of Cartesianism grew up, indeed, evolved, together. A close relative of Descartes' ontological internalism is epistemic internalism, a view which has its modern

roots in Descartes. The central idea of epistemic internalism is that the difference between true belief and knowledge consists in some form of justification and, crucially, that justification consists in factors that are, in some sense, internal to the subject of the belief. The relevant notion of internality, however, is fundamentally epistemic. The activities of my heart, lungs, and liver are activities internal to me, but clearly these are not candidates for transformers of true belief into knowledge. Whatever transforms true belief into knowledge is, according to internalism, something of which the believing subject can be aware, something to which the subject has epistemic access. I can, however, be aware of many things, including whether or not it is presently raining, and the week's activity on the New York stock exchange. But this is not access of the relevant sort. Rather, epistemic internalism claims that justification consists in some sort of special access. According to Descartes, for example, the special access consists in the fact that the thinking subject can determine with certainty whether a belief has justification, and, consequently, whether a belief qualifies as knowledge. And, according to Chisholm (1966), a recent internalist descendant of Descartes, whether a belief has justification is something that can be determined by reflection alone. So, the central idea of epistemic internalism is that the factors that make a true belief justified, and, consequently, that transform a true belief into knowledge, are properties to which the believer has a special sort of epistemic access.

All these theses would be, if not straightforwardly undermined, then at least significantly threatened by the rejection of the Cartesian conception of the mind. If the mind is not self-contained in the relevant way, if the world enters into the very constitution of the mind so that the very identity of mental states involves something external, then it is not clear that this leaves any room for the special sort of epistemic access required by the internalist tradition. For the identity of the contents of one's mind would now, in part, consist of items to which one had no special epistemic access, and this threatens the epistemic relation one bears to the contents of one's mind as a whole. In this way, Descartes' ontological internalism is bound up with, supports and is supported by, his epistemic internalism.

Indeed, the content of internalism extends even beyond the ontological and epistemological spheres into the moral domain. This is because epistemic internalism is closely connected with the deonto-

logical notion of epistemic responsibility. The justification, or lack thereof, of my beliefs, is something for which I am responsible, something for which I can be praised or blamed. I may be victimized by a malevolent demon, I may be a brain in a vat, and so all or most of my beliefs may, in fact, be false. The truth of my beliefs depends on external factors and so is something beyond my control. But the justification of my beliefs depends only on internal factors, indeed, internal factors to which I have a special kind of access, and this is something that does lie within my control. Even though I may be hopelessly deceived about the truth of my beliefs, I can still do my epistemic duty with regard to their justification. Accordingly, the justification of my beliefs is something for which I can legitimately be praised or blamed; it is something for which I can be morally assessed. Thus, according to Descartes, epistemic justification is a form of deontological justification. If I do not have certainty but believe anyway, then I do not escape the blame of misusing my freedom. Ontological and epistemic forms of internalism, then, are also closely bound up with a certain view of the moral nature and responsibilities of human beings.

To speak of the Cartesian conception, therefore, is to speak of not just a single view of the mind, but of an array of interwoven views, each lending support to the others, and each being supported by the others. The strength of the Cartesian picture lies not merely in the strength of the individual theses – ontic, epistemic, ethical – that constitute it, but also, and perhaps even more importantly, in its scope. The Cartesian picture provides us with a sweeping and comprehensive vision of the nature of human beings. And the strength of the individual components of this vision derives, in an important sense, from the strength of the vision as a whole.

This book, in one clear sense, seeks to undermine the Cartesian picture of human beings. However, its primary concern is not with the picture as such, but with the influence it has exerted on subsequent theorizing about the mind. Moreover, the principal focus of the book will be restricted to a sub-set of mental phenomena that have proved particularly central to twentieth-century concerns. These phenomena comprise what are known as cognitive processes: processes such as perceiving, remembering, and reasoning whereby an organism gains and uses information about its environment. The Cartesian picture has bequeathed us a conception of such processes

whose essence can be distilled into the following two principles. One is an ontological claim about the nature of cognitive processes, the other an epistemological corollary about how these processes are best studied or understood.

The Ontological Claim: Mental states and processes are located exclusively inside the skin of cognizing organisms.

The Epistemological Claim: It is possible to understand the nature of mental states and processes by focusing exclusively on what is occurring inside the skin of cognizing organisms.

These two assumptions constitute that particular version of the internalist picture of the mind with which this book is concerned. Cognitive processes are *essentially* internal items. They may stand in various relations to events, states, and processes occurring outside the skin of cognizers, and these external items may play an important, even essential, role in the facilitation or satisfaction of the internal processes themselves. Nevertheless, cognitive processes are, in essence, internal items. This internalist picture of cognition, I shall try to show, has greatly distorted our conception of what is required of us by the projects of psychotectonics and psychosemantics. And the principal task of this book is to unseat this conception of cognition, and outline, in broad strokes, the ramifications of this for the two projects.

1.2 PICTURE AS MYTHOLOGY

The task of this book is to unseat not a particular philosophical or psychological theory, but a certain pre-theoretical conception or picture of the mind. This picture is prior to theory in that it is what guides theory construction and thus lends coherence and unity to the experimental practices and procedures judged relevant to the conformation or falsification of particular theories within its domain. The notion of a pre-theoretical picture, then, corresponds largely to what Kuhn (1970) has, famously, labelled a *paradigm*.

In trying to unseat a pre-theoretical picture of a certain domain of inquiry, one's options are usually fairly restricted. One might try to attack the picture on grounds of internal incoherence. Many pictures, perhaps most, however, are not internally incoherent, and, certainly, I

would not want to suggest that the internalist picture of the mind suffers from this sort of defect. One might try to attack the picture on grounds of empirical inadequacy. As Wittgenstein, Kuhn, and others have pointed out, however, this sort of approach is unlikely to work. The problem is not so much that a pre-theoretical picture is at a further remove from the evidence, hence at a further remove from the possibility of experimental confirmation or falsification, than the theories predicated upon it (although this certainly might be a problem). The problem, rather, is that the very descriptions of the evidence employed for the purposes of experimental testing are based on, and thus presuppose, this pre-theoretical picture. It is very difficult to see how a pre-theoretical picture of a domain could be unseated solely on the basis of evidence whose very identification and conceptualization as evidence, presupposes the picture itself.

Kuhn, as is well known, sees the unseating of a pre-theoretical picture or paradigm as a matter of piecemeal accretions of problems unsolved within the framework of the picture gradually inducing a *crisis* within the picture itself. Such a crisis is likely to come to a head only when there is a competitor to the picture, an alternative paradigm that can play the same sort of role as the original. As Kuhn points out, however, the new picture, initially, is likely to be partial, restricted, and susceptible to all sorts of apparent refutations. This is because the evidence thought relevant to its truth or falsity is still conceptualized in terms of the old paradigm. It is only when the new paradigm becomes accepted that the relevant evidence can gradually be reconceptualized, and then the scope and coherence of the new paradigm can be progressively enhanced. Arguably, such a situation may today be occurring in cognitive science. Arguably, it may not. And, in any event, I would not want to predicate any argument upon such an essentially contestable claim.

The inspiration for the method I propose to adopt in this book derives from Wittgenstein rather than Kuhn. A pre-theoretical picture in the sense described above has the status of what Wittgenstein calls a *mythology*. To call something a mythology in this sense is not to cast, at least not directly, any aspersions at its truth or validity, although it may mean that, for it, questions of truth and validity do not arise. What Wittgenstein is getting at when he calls something a mythology is that it plays a certain role in organizing experience. More precisely, a mythology has the role of legislating

that *this is how things must be*. One of the principal tasks of Wittgenstein's later philosophy is to show how certain pre-theoretical conceptions of the mind and its contents are mythologies in precisely this sense. And one of the principal methods employed by Wittgenstein in this context is to undermine a mythology by showing that we do not, in fact, have to think of things in the way the mythology tells us we do. Other ways of thinking about the mind and its contents, for example, are possible.

A good example of Wittgenstein's idea of mythology is to be found in his attitude toward Freud's concept of the unconscious. In his *Cambridge Lectures 1932–35*, Wittgenstein writes:

What Freud says about the subconscious sounds like science, but in fact is just a *means of representation*. New regions of the soul have not been discovered, as his writings suggest. The display of elements of a dream, for example, a hat (which may mean practically anything) is a display of similes. As in aesthetics, things are placed side by side so as to exhibit certain features. (1979:40)

Wittgenstein is quite willing to allow that Freud has discovered certain psychological reactions of a hitherto unknown sort, but the apparatus he invokes to explain these is not a theory but simply a *means of representation*. That is, he has simply imposed, as Wittgenstein would say, a system of notation which allows him to redescribe these psychological reactions in these terms. Psychoanalysis, while presenting itself as an experimental discipline, does not, in fact, satisfy any of the conditions necessary to a discipline of this kind.

The comparison with aesthetics is indicative of Wittgenstein's attitude towards the unconscious. Consider the difference between the role of analogy ('simile') in fields like aesthetics and its role in the empirical sciences. An analogy of the first type might consist, for example, in comparing architecture with a language and, then, attempting to identify the vocabulary and grammar of this language. This type of analogy, however, does not generate hypotheses that can be tested in experiments, nor does it produce a theory that can be used to predict events. Thus, whatever understanding is occasioned by the use of such analogies is not the result of imparting new information, nor does it lead to new empirical discoveries. Furthermore, such understanding does not lead to the asking of fresh questions that can be answered by further empirical research. The analogy, rather, func-

tions by making formal connections between architectural features and linguistic ones. The analogy yields new forms of comparison, changing our understanding of buildings and altering the way we look at things. We can then, for example, describe a piece of architecture in a new way: as making sense or not; as being rhetorical or bombastic; as being witty or ambiguous, and so on (Hacker 1987).

According to Wittgenstein, what Freud does is essentially offer us good analogies. But these analogies are of the sort employed by art historians and art critics, not the sort used by physicists. As Hacker (1987:487) puts it, the analogies employed by Freud are not *model-generating* but simply *aspect-seeing*. In presenting the hypothesis of the unconscious as a scientific hypothesis, then, Freud has misunderstood the type of illumination the hypothesis provides; he has misunderstood the status of the type of inquiry in which he is engaged.

In presenting the aspect-seeing analogy of the unconscious as if it were a model-generating scientific hypothesis, Freud has, in effect, introduced a new mythology into the study of the mind. For our purposes, the important feature of a mythology, in Wittgenstein's sense, is the role it plays in delimiting the acceptable forms of description and explanation applicable in the domain in which the mythology is operative. In the *Remarks on Colour*, Wittgenstein cites the Freudian idea of dreaming as disguised wish-fulfilment as an example of what he calls a *primary phenomenon*: 'The primary phenomenon is, e.g., what Freud thought he recognized in simple wish-fulfilment dreams. The primary phenomenon is a preconceived idea that takes possession of us' (no. 230). The primary phenomenon acts as a model or prototype that is to be used in the description of phenomena. And such a prototype, for Wittgenstein, 'characterizes the whole discussion and determines its form' (Wittgenstein 1978:14). In Freud, the model of dreams as disguised wish-fulfilment, a model whose legitimacy depends on the postulation of the unconscious, is presented as a scientific hypothesis, as a discovery about the real nature of dreaming. In fact, however, at least according to Wittgenstein, the model is something that determines the form of description and explanation of all phenomena considered relevant. It, therefore, applies to dreams not because it has been demonstrated by a scientific investigation of different kinds of dreams but because it has been granted a privileged place in discussion. The result is not simply that all evidence can now be described so as to confirm Freud's hypothesis, evidence that might

prima facie just as easily disconfirm it. More importantly, with respect to the phenomena of interest, the *grammar* of any possible reason or explanation has been determined, and any competing reason or explanation that falls outside certain parameters cannot be accepted for consideration as the explanation or the reason of that phenomenon.

A mythology, for Wittgenstein, is essentially a pre-theoretical picture which has the function of delimiting, in advance, the form or grammar of all possible descriptions and explanations of the phenomena that fall within its scope. Thus, it has the function of making us see, or understand, the phenomena in a certain way rather than another. It makes us think *this is how things must be*, and initiate our explicit theorizing from this starting-point. Its efficacy in this role depends, to a considerable extent, on us, as in the case of Freud, not recognizing this mythology for what it is; on our mistaking it for something else, a scientific hypothesis, for example.

The central argument of this book is that our thinking about the mind has been guided by a mythology in roughly the sense described above; the myth is that of internalism. The goal of this book is to enable us to see this idea for the mythology that it is, and to thereby undermine the influence it has on our explicit theorizing about the mind and its contents.

1.3 SUBVERTING MYTHOLOGY

The aim of this book, then, is to unseat a certain pre-theoretical picture of the mind. The term *unseat* is not arbitrarily chosen. The aim of the book is *not* to *refute* this picture, but to unseat it. The aim, that is, is not to show that the picture is internally inconsistent, or lacking on empirical grounds. It may be either or both of those things, but I am not claiming that it is, and certainly not basing any argument on assumption that it is. The aim of the book is to unseat the picture in the Wittgensteinian sense of showing that this is not how we have to think about cognition. I shall argue that there is no theoretically respectable reason for separating the mind off from the world in the way the internalist picture tells us we should. There is, in other words, no theoretically respectable reason for thinking of cognitive processes as purely and exclusively internal items. And to say there is no theoretically respectable reason, here, simply means that there is no reason

that can be derived from psychological theory as such. The parsing of the realm of cognition into, on the one hand, cognitive processes that are conceived of as purely internal items and, on the other, external causes, stimuli, or cues of these internal items is not something that is demanded by our theorizing about the mind, but an optional extra. It is a pre-theoretical picture we use to interpret our explicit theorizing, not something mandated by that theorizing. It is, in short, a mythology.

The aim of this book is to subvert or unseat this pre-theoretical picture. The strategy is to show that this is not how we have to think about the mind by outlining an alternative picture, and accumulating various sorts of evidence in favour of this picture. Thus, while I shall offer an alternative pre-theoretical picture of the mind, the rationale for offering this alternative is, to a considerable extent, negative: its primary aim is to unseat another picture, and only secondarily does it present a positive basis for new or additional theorizing. The project of psychotectonics, as developed in Part I of this book, firstly, acts as a palliative or antidote to a conception of the mind that has exerted a vice-like grip on our theorizing, and, secondly, indicates, in admittedly broad strokes, how our theorizing might change should we embrace this alternative conception.

Another feature of Wittgenstein's method is relevant to the approach adopted by this book. One of the more surprising features of Wittgenstein's approach is that he frequently seems to be attacking straw men. That is, he often appears to be attacking extremely simplified, indeed oversimplified, versions of the views he challenges. His attack on the Augustinian account of language acquisition in the first part of the *Philosophical Investigations* is a good example of this. Augustine's actual position is a lot more subtle and complex than one would suspect from a cursory reading of Wittgenstein, incorporating, in fact, many features with which Wittgenstein would, presumably, be in agreement. Moreover, more recent versions of the Augustinian account – into which category the *Tractatus Logico-Philosophicus* broadly falls – are certainly much more subtle, containing complexities Wittgenstein does not even begin to address. It may legitimately be wondered why Wittgenstein ignores the complexity of Augustine's view, or developments of the view, and attacks a version of it that is so simplified as to amount almost to parody. The answer is that Wittgenstein is concerned with unseating not any particular

theory of meaning but, rather, the *picture* that lies behind it. And to get at the picture, one must strip away the details. The picture, for Wittgenstein, is a type of mythology in that it leads us to think that things must be a certain way. And the goal of Wittgenstein's method is simply to show that we do not have to think of things like this. Things do not, in fact, have to be this way at all. One unseats the picture by loosening its grip on us. And one way of doing this is to describe, in broad outline, another way of looking at things. And, for this task, the specific details of the various theoretical developments of the picture are irrelevant.

This procedure is also, to a considerable extent, adopted in this book. This is particularly true of Part I. The task of this part is to unseat the internalist picture of the mind described above. The strategy is to loosen the grip this picture has on us. And the method is to outline and defend an alternative way of looking at mental processes. Adopting this strategy leads me to describe various theoretical developments of the internalist picture in ways that are oversimplified. Nowhere, for example, is this more true than in the discussion of David Marr's theory of vision in chapter 5. Marr's theory is, in fact, a lot more subtle, complex, and indeed plausible, than would be apparent from a cursory reading of chapter 5. Indeed, Marr incorporates many ideas that are compatible with, even congenial to, the alternative environmentalist picture outlined in this part of the book (I am here thinking particularly of the theoretical role Marr gives to the notion of *assumptions* about the environment). There is a clear sense, I think, in which Marr's account is not straightforwardly in opposition to the Gibsonian account outlined in that chapter. A more accurate way of thinking about the relation between the two theories is, perhaps, to regard them as taking up positions on a spectrum. The poles of the spectrum are constituted by the internalist picture described above, and the alternative environmentalist picture outlined in the course of Part I. The difference between Marr's theory and that of Gibson, then, would be explained in terms of their relative positions on the spectrum. Marr, I think it is safe to say, is closer to the internalist end of the spectrum than Gibson. However, the precise position occupied by Marr, and indeed by Gibson, on this spectrum, is a question for the sort of detailed textual interpretation that is beyond the scope of this book. If you think, for example, that Marr ought to be located considerably more towards the environmentalist

end of the spectrum than I have given him credit for, that is fine with me. My concern with Marr is restricted solely to his (as I see it) role as representative of the internalist position with respect to visual perception. My concern throughout Part I of the book, then, is not with the details of the specific theoretical developments of the internalist picture; it is with the picture itself. And this concern will, at certain points, lead me to overlook many of the details of its theoretical articulation, details which, in other contexts, and for other purposes, would be extremely important.

The central purpose of this book, then, is to unseat a certain pre-theoretical picture of the mind by loosening the grip it has on our thinking about the mind and mental processes. It attempts to do this by showing not that this picture is false or incoherent, but by showing that it is not necessary for our thinking about the mind; that we do not have to think about the mind and its contents in the way the picture requires. Failure to understand the strategy will almost certainly result in failure to understand the arguments. For example, in chapters 7 and 8 there is an extended discussion of connectionist approaches to the modelling of cognition in relation to the traditional symbolic approaches. This discussion is presented as one particular case in which the internalist picture has exerted a rather baleful influence on our theorizing. And, correspondingly, liberation from this picture can, I think, help us to understand more fully the merits of connectionism. However, nowhere in this discussion will one find the conclusion: connectionism is right and the symbolic approach is wrong. My concern in these chapters is not with connectionism or the symbolic approach as such, but with the pre-theoretical pictures that these programmes articulate. And, once again, the same tendency to oversimplification conspicuous in the discussion of Marr will also be evident. The purpose of the book is not to show that the internalist picture is false or incoherent. Therefore, neither is it to show that the theoretical articulations of this picture – as, I will argue, are the symbolic approach and associated language of thought hypothesis – are false or incoherent. The purpose, rather, is to loosen the grip that the dominant internalist picture has on us by showing that we do not have to think of things in this way. And this, I shall try to show, can shed new light on the dispute between connectionist and symbolic approaches. The aim of chapters 7 and 8, then, is not to show that the symbolic approach to understanding cognition, and the associated

language of thought hypothesis, are false. The aim is to undermine them by removing, at least in part, their motivation.

To some extent already, the grip of the internalist picture is relaxing. In an excellent study, Andy Clark (1997) has shown how the internalist picture is, at best, dubiously compatible with recent development in artificial intelligence, in particular, development in the field of robotics and artificial life. I accept Clark's arguments, and the arguments to be developed in the following pages are, in an important sense, of a piece with his. However, there are differences of strategy. Firstly, while Clark focuses primarily on very recent work in the rather new fields of robotics and artificial intelligence, I shall focus primarily on older work in more traditional areas of psychology. This difference in strategy stems largely from my goal of unseating a certain picture of the mind, rather than of demonstrating its inadequacy. The goal of this book is as much persuasion as anything else; the goal is to persuade us that we do not have to think about the mind in this way, that other ways of looking at the mind are possible, indeed viable. And, a useful way of doing this is by showing that this is something we have always really known anyway, or, at least, something that is implicit in a quite familiar body of knowledge.

Paul Feyerabend (1975) has argued, plausibly, that this was the sort of strategy adopted by Galileo in his attempt to develop an alternative to Aristotelian cosmology. One of the key concepts required for this development was that of relative motion. The postulation of a moving earth, for example, was apparently quite clearly incompatible with certain uncontroversial facts of experience; for example, that rocks dropped off the top of towers hit the ground at the bottom of the tower, and not some distance away as might be expected if the earth were really moving. To undermine these apparently watertight experiential refutations, Galileo was obliged to develop the concept of relative motion. What is interesting is that his strategy here was to try and show that the concept of relative motion – a concept quite alien to Aristotelian cosmology – was something with which we were all quite familiar anyway. Feyerabend compares this to Plato's idea of *anamnesis*: the explicit remembering of something we have always implicitly known. This, at least to some extent, is the strategy that will be adopted in Part I of this book. The strategy, therefore, is to focus on work in psychology that is not particularly recent and is more or less traditional (with the possible exception of Gibson). And in this work,

it will be argued, we find no theoretically respectable reason for embracing internalist mythology.

In Part II, the focus switches from psychotectonics to psycho-semantics. This is for two reasons. Firstly, the alternative picture of the mind developed in Part I relies, as we will see, quite heavily on the notion of mental representation. Thus, the completeness of the alternative anti-internalist, or as I shall call it, *environmentalist*, picture of cognitive processes requires an account of mental representation. If we can assume that the extent to which the environmentalist picture can be seen as a legitimate alternative to the internalist one depends in part on its completeness, then the success of the project of Part I requires the development of an account of mental representation. This is the project of Part II.

Just as importantly, however, it will be argued in Part II that recent discussions of the notion 'mental representation' have suffered considerably from a tacit adherence to the internalist mythology that Part I tries to undermine. Thus, it will be argued, replacing the internalist picture of cognition with the environmentalist account developed in Part I allows us to solve, or dissolve, some of the most troubling objections to naturalistic accounts of representation. Thus, the arguments of Part II serve a dual function. On the one hand, they complete the arguments begun in Part I. On the other hand, they give one important example of why those arguments were necessary.